TITLE PAGE

Selected Paper #174502 prepared for presentation at the American Agricultural Economics Association Annual Meeting, Portland, OR, July 29-August 1, 2007

Title of Paper: Identifying ENSO Phase Impacts on Area Yield Insurance Rates: An Application of Non-Parametric Analysis

Authors:
Denis A. Nadolnyak
Department of Agricultural Economics and Rural Sociology
312 Comer Hall
Auburn University
Auburn, AL 36849
Phone: (334)844-5630
E-mail: nadolda@auburn.edu

James L. Novak
Professor, Department of Agricultural Economics and Rural Sociology
Auburn University
304 Comer Hall
Auburn, AL 36849
Phone: (334)844-3512
E-mail: NOVAKJL@auburn.edu

Joel O. Paz
Department of Biological and Agricultural Engineering
University of Georgia
1109 Experiment St.
Griffin, GA 30223-1797
Phone 770-228-7399
E-mail: jpaz@uga.edu

JEL codes: Q140, C220, G220.

Identifying ENSO Phase Impacts on Area Yield Insurance Rates: An Application of Non-Parametric Analysis

Authors

Denis A. Nadolnyak, Department of Agricultural Economics, Auburn University
James L. Novak, Department of Agricultural Economics, Auburn University
Joel O. Paz, Department of Biological and Agricultural Engineering, University of Georgia

Selected Paper prepared for presentation at the American Agricultural Economics Association Annual Meeting, Portland, OR, July 29-August 1, 2007

INTRODUCTION

The ElNino Southern Oscillation (ENSO) phases have distinct impacts on the climate in the Southeastern United States. It is reasonable to expect that they also affect crop yields (Hansen, Hodges, and Jones). If this is the case, it is important to know whether the effects are significant enough to be accommodated in agricultural decision making.

The ElNino Southern Oscillation (ENSO) is an atmospheric phenomenon observed with irregular periodicity which is believed to affect global climate. The phases are determined by warming or cooling of the ocean surface in the western Pacific ocean, which changes the trade wind patterns and, subsequently, global weather for the duration of the phase and beyond (see http://meted.ucar.edu/climate/enso/index.htm for the basics of the process). The effects of ENSO phases on climate in different geographical regions are complex and are not discussed here.
In the Southeastern United States, more or less distinct ENSO dependent climate patterns have been observed. In addition, the likelihood of a severe freeze is much greater during a neutral phase than during either an El Nino or a La Nina event. The table below provides a summary of the effects on temperature and precipitation across the Southeast.

El Nino/La Nina Impacts Across the Southeast U.S.					
Phase	Region	Seasons			
		Oct-Dec	Jan-Mar	Apr-Jun	Jul-Sep
El Nino	Peninsular Florida	Wet & cool	Very wet & cool	Slightly dry	Slightly dry to no impact
	Tri-State Region	Wet	Wet	Slightly wet	No impact
	Western Florida Panhandle	No impact	Wet	Slightly dry	No impact
	Central and North Ala. & Ga.	No impact	No impact	No impact	Slightly dry
La Nina	Peninsular Florida	Dry & slightly warm	Very dry & warm	Slightly wet	Slightly cool
	Tri-State Region	Slighly dry	Dry	Dry	No impact
	Western Florida Panhandle	Slightly dry	Dry	Dry	No impact
	Central and North Ala. & Ga.	Dry	Dry in the south, wet in NW Ala.	No impact	Wet in NW Ala.
Neutral	All Regions	No impact	No impact	No impact	No impact

Source: http://www.coastalclimate.org/climate/seimpacts.php

In this paper, we report results of non-parametric analysis of peanut, corn, and cotton yield distributions by the ENSO phases in the Southwestern Georgia, Northwestern Florida, and Southern and Northwestern Alabama. For comparison and validation purposes, historical yield data are complemented by a set of simulated peanut yields generated using daily weather, soil, and management practices data for Colquitt county, Georgia.

The focus of the analysis is on establishing ENSO-dependent differences in the yield distributions and on evaluating their implications for area yield crop insurance the expected losses for which are calculated using county average yield series. The hypothesis is that different climate conditions during ENSO cycles translate into different yield distributions, which is justified by the observed South-Eastern climate differences and previous research showing the importance ENSO cycles for optimal planting dates.

For estimating yield distributions, we use the non-parametric technique of kernel density estimation as it appears more suited for the purpose than the parametric methods in that it accommodates skewness, bi-modality, and other peculiarities observed in crop yield densities.

The results of kernel density estimates of historical county yield data show consistent patterns in actuarially fair insurance rate schedules grouped by ENSO phases and geographical areas. In particular, corn yield insurance premiums appear to be the most dependent on the ENSO phases and are the highest, regardless of coverage, during ElNino and the lowest during LaNina. Cotton premiums are the highest for ElNino (except for North-Central AL) and the lowest during LaNina. Peanut premiums are higher during Neutral years and lower during LaNina (except of Southern Alabama). The results appear to be robust to the transformations used to make the yield series stationary. While these dependencies do not necessarily correspond to the precipitation and solar radiation characteristics of the corresponding ENSO cycles in the Southeastern US, drawing direct analogies with yield variability is premature as many less documented factors, like the spacing of sunny and rainy days, may be just as important.

The validity of these findings is reinforced by comparison of the empirical and simulated peanut yield distributions. The comparisons show that the distributions are similar in many ways and that the dissimilarities can be explained by known factors. These findings should be more relevant for the area yield insurance as opposed to the APH arrangements as the yield data used in designing contracts for the former reflects the systemic risk more influenced by climate than by the farm-level, basis risk factors accommodated in the APH plans.

The rest of the paper is structured as follows. The Methodology section explains the density estimation techniques and insurance rate calculation. The Simulated Yields section describes the peanut yield simulation process and the resulting yield distribution by ENSO phases. The Historical County Yields section describes methodologies used to make the series stationary and compares the simulated and empirical distributions. The Insurance Rates section presents the results of area yield insurance premium calculations for corn, cotton, and peanuts in Georgia, Florida, and Alabama.

Methodology

Non-Parametric Density Estimation

Conventional parametric approaches to insurance analysis assume known functional forms for yield distributions. The most commonly assumed density is normal, which is justified by the Central Limit Theorem. However, empirical yield data do not always conform to theoretical priors due to a number of physical and biological attributes of plant growth. In particular, bi-modality and skewness of yield distributions are often observed. Non-parametric density estimation accommodates these and other distributional idiosyncrasies.

The simplest way to estimate non-parametric density is to use a histogram. This, however, poses the problem of discontinuity and requires large samples. Smoothing the density between observations utilizes a kernel function, an estimator of local density around a datum. Each observation is surrounded by a symmetric weighting (probability density) function K satisfying:

$$\int_{-\infty}^{+\infty} K(t)dt = 1$$

The kernel estimator for n IID observations of Y is given by

$$\frac{1}{nh}\sum_{1}^{n} K\left(\frac{y-Y_i}{h}\right)$$

where h is the bandwidth parameter assigning the weight to neighboring observations (the amount of smoothing), usually chosen on the basis of minimizing the mean integrated square error.

The local density estimates overlap (to the degree of the kernel width), so that each kernel density depends not only on its own but also adjacent observations. The resulting density estimation is continuous and is a more adequate representation of true yield distributions than parametric densities. Kernel density estimates can also be compared to other (non)parametric densities using Kolmogorov-Smirnov tests. There is a range of possible kernel functions and kernel widths, the latter arguably being a more important specification. For this analysis, we chose Gaussian kernels as more commonly used in economics and the kernel width according to the Silverman's "rule of thumb" as optimal for the normal distribution family. However, different specifications did not dramatically change the results. For a more thorough discussion of kernel density estimation, see Li and Racine.

Non-Stationarity of the Yield Series

In order to properly single out the impact of the ENSO phases on yield distributions, it is necessary to remove trends, autoregressive effects, heteroscedasticity, structural breaks, and other systematic influences affecting the distributions. Many of these are expected in the yield time series mainly due to technological changes (trends and heteroscedasticity) and drought/moisture effect persistence (auto-regression and moving average). A variety of tests and methodologies are available for dealing with these problems, the results of which are reported in the Non-Stationarity and Detrending sub-section of the paper.

Crop Area Yield Insurance

Area yield, or group risk (GRP), insurance was chosen for the analysis because, unlike individual (farm-level) crop insurance based on actual production history, area yield insurance contracts are based on yield forecasts. Besides, county yields comprise often offsetting farm level influences and thus are more likely to depend on climate variability as a universal impact. The major disadvantage of the county yield data is that it does not distinguish between irrigated and non-irrigated yields.

A basic GRP contract insures a certain percentage of the expected yield (the coverage usually ranging from 70% to 95%), the indemnity being equal to the difference between the covered percentage of the expected yield and the actual below coverage yield. The expected loss for the insurer is thus a probability weighted sum of indemnities for all possible yield realizations:

$$Expected\ Loss = \int_{\min\_yield}^{E[yield]*coverage} f(Y)*(\overline{Y}*coverage - Y)dY \qquad ,$$

where $\overline{Y}$ is the expected yield and *f(Y)* is the yield distribution density. The actuarially fair premium, i.e., the risk-neutral insurer's break-even point, is a ratio of the expected loss and $\overline{Y}*coverage$. While the actual rates are different because of the loading factors and subsidies, the fair premium is at least a component and provides a good ceteris paribus reference point. Actuarial fairness is also important for the risk-averse producer (the insured) as it corresponds to a competitive industry outcome and is thus welfare improving.

Obviously, the accuracy of insurance premiums depends on the area below the yield's pdf and to the left of the coverage. Thus, the exact shape of the underlying distribution and/or the ability of its estimate to accommodate skewness, bi-modality, and other idiosyncrasies observed in crop yield data are important for establishing actuarial fairness and contract efficiency. The figure below illustrates: a right skewed distribution with the same mean may result in a smaller expected loss than the one for an otherwise identical but symmetric density.

Figure: expected losses and the density shape

a) symmetric pdf b) right skewed pdf

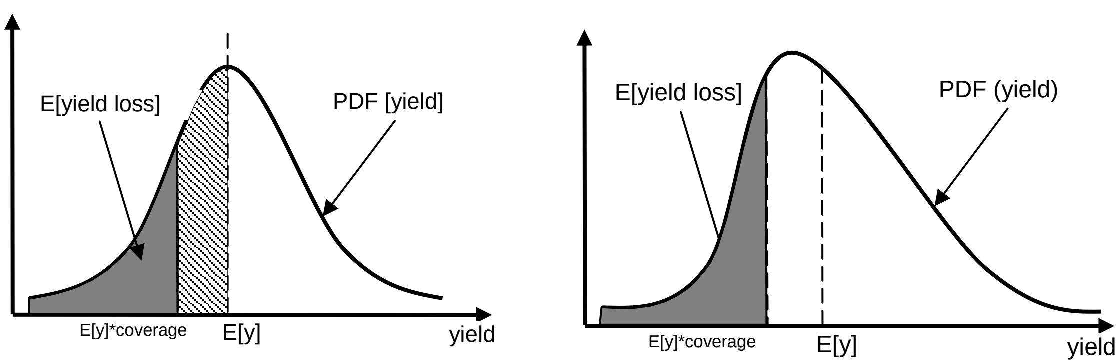

In the analysis, we use the trapezoid rule for integrating the empirically estimated pdf's. For a sufficiently large number of grid points generated by kernel density estimation, it provides accurate enough estimates.

Simulated Yield Data

Crop Simulation Methodology

The simulated peanut yield dataset was generated using the Cropping System Model (CSM)-CROPGRO-Peanut model (Boote, Jones, and Hoogenboom, 1998; Jones et al., 2003) and made available to us courtesy of Joel Paz (UGA) and Clyde Fraisse (UF). The CSM-CROPGRO-Peanut model, which is part of Decision Support System for Agrotechnology Transfer (DSSAT) Version 4.0 (Hoogenboom et al., 2004), is a process-based model that simulates crop growth and development. Long-term historical weather data (1900-2004) were obtained from the National Weather Service (NWS) Cooperative Observer Program (COOP) network and compiled by the Center for Oceanic-Atmospheric Prediction Studies (COAPS), through the South-Eastern Climate Consortium (SECC). The weather variables include daily maximum and minimum temperatures and precipitation. A solar radiation generator, WGENR, with adjustment factors obtained for the southeastern USA, was used to generate daily solar radiation data.

Georgia Green was selected as the representative variety for Georgia and other southeastern states. The soil profile data of three representative soils for each Colquitt county, GA, were obtained from the soil characterization database of the USDA National Resource Conservation Service. Nine planting dates (April 16 – June 12) represented all possible planting dates at weekly intervals. Peanut yields were simulated assuming no irrigation.

Regardless of their complexity and accommodation of biological and physical processes, the crop simulation models are deterministic. Therefore, whatever randomness in simulated yields is observed for same plots and management practices comes from random weather realizations. In this way, the simulated data is analogous to a controlled experiment. At the same time, it is nearly impossible to translate weather variability, expressed in so many ways, into yield variability through the model mechanics. For instance, cumulative measures of precipitation and solar radiation may not be correlated with yields if the weather patterns are different, as evidenced by a comparison of the effect on plant growth of a week with four rainy days each followed by a sunny one with a week in which it rains four days in a row (the first one is likely to be more favorable for growth). Thus, we do not try to deliberately draw parallels between climate indexes and our findings. Instead, we independently estimate the distributions of the simulated yields without forming any a priori expectations based on climate research.

The simulated annual data covers the period from 1911 to 2003 and assumes modern “best” management practices. This time period covers 14 ElNino, 17 LaNina, and 39 Neutral years. This is barely enough for distribution analysis, but the actual daily weather observation records do not go back much further. The records are from a weather station in Colquitt County in Southwestern Georgia, located in the heart of the Southeastern peanut producing region.

The nine simulated planting dates and three soil types make for 2511 observations. The three soil types assumed are Tifton Loamy Sand, Cowarts Loamy Sand, and Troup Sand, the first being the most prevalent in the county (NRCS). As the differences in yields between the soil types were negligible and because peanuts are planted on all these soils, we did not distinguish between the soil types in most of the analysis.

Simulated Yields Analysis

The table below provides some basic parameters of the distribution of the simulated peanut yields by the ENSO phase.

Simulated peanut yield distribution, average of planting dates and soils

enso	mean	sd	skewness	kurtosis	*min*	*max*
ElNino	2329.25	1003.81	0.43	2.92	*357*	*4902*
LaNina	2588.70	1282.87	0.31	2.03	*337*	*5519*
Neutral	2298.67	1179.43	0.85	3.29	*454*	*6297*
All years	2373.53	1177.45	0.66	2.85	*337*	*6297*

LaNina yields have the highest mean, which is associated with higher variance. The skewnesss of the simulated yield distribution during neutral years is fairly high (>0.5), while during the ElNino and LaNina phases it is below moderate. Higher kurtosis is normally interpreted as greater “peakedness”, which means more of the variance is due to infrequent extreme deviations, as opposed to frequent modestly-sized deviations. However, these moments convey more information when applied to parametric distributions.

The statistics for different planting dates and soil types show that the highest yields come from planting in between May 8-22 and Tifton Loamy Sand. However, the *relative* values of the distribution statistics are preserved throughout. As evaluating soil productivity and management practices is beyond the scope of this paper, we report such differences only in cases where they affect the distributional differences.

Assuming the simulated data represents a controlled experiment in that the yields do not depend on unknown stochastic influences, the important question is whether the yields depend on the ENSO phase. The results of t-tests of mean equality are reported in a table below.

Mean equality tests, simulated peanut yields

Ha: diff < 0	Ha: diff != 0	Ha: diff > 0
H_0: mean(yield, EL) – mean(yield, LA) = 0		
Pr(T < t) = 0.0001	Pr(|T| > |t|) = 0.0002	Pr(T > t) = 0.9999
H_0: mean(yield, EL) – mean(yield, Neu) = 0		
Pr(T < t) = 0.6992	Pr(|T| > |t|) = 0.6016	Pr(T > t) = 0.3008
H_0: mean(yield, LA) – mean(yield, Neu) = 0		
Pr(T < t) = 1.0000	Pr(|T| > |t|) = 0.0000	Pr(T > t) = 0.0000

The tests confirm that LaNina year yields are the highest in comparison to Neutral and ElNino, while there is no statistical difference between Neutral and ElNino. It has been suggested that the effects of the ENSO phases can carry over to the next year, but we did not find indications of such influences in the simulated data.

The significant difference in average yields depending on the ENSO phase is an interesting find, especially considering the fact that little ENSO climate impact has been found in south Georgia, south Alabama, and north-western Florida. One reason for this discrepancy might lie in the details of the crop growth functions and their dependence not on the average temperature or precipitation, but on the finer details of climate, like the spacing of rainy and sunny days.

As the moments do not describe empirical distributions completely, we proceed by reporting the differences between the distributions. As common procedures for testing equality of variances rely on distributional assumptions which might not hold for the yield data, we use the non-parametric Kolmogorov-Smirnov test. The two sample test is based on the maximum absolute difference (D) between the CDFs for two continuous random variables. Unlike conventional statistical tests, this is a non-parametric test that does not require the variables to be normally distributed. The null hypothesis for the Kolmogorov-Smirnov test is that there is no difference in the CDFs between two groups. The largest observed difference between the two CDFs being examined was compared to the critical value of D at the 5 percent level of significance to determine if there is a statistically significant difference between the curves. The table below reports the results.

Kolmogorov-Smirnov tests for distribution differences

H_0: f(EL) = f(LA)			
Smaller group	D	P-value	Corrected
EL	0.1702	0.000	
LA	-0.0595	0.142	
Combined K-S	0.1702	0.000	0.000
H_0: f(EL) = f(Neutral)			
Smaller group	D	P-value	Corrected
EL	0.0471	0.189	
Neutral	-0.1188	0.000	
Combined K-S	0.1188	0.000	0.000
H_0: f(LA) = f(Neutral)			
Smaller group	D	P-value	Corrected
LA	0.0164	0.799	
Neutral	-0.1333	0.000	
Combined K-S	0.1333	0.000	0.000

The combined K-S statistics clearly show that the yield distributions in El, La, and Neutral years are significantly different from each other (the differences are similar across planting dates and soil types).

The observed differences in the distributions have immediate implications for insurance design, particularly for the group risk plans where rates for incomplete (less than 100%) coverage are calculated on the basis of distributional parameters. Such calculations may assume normality or use exponential smoothing but without regard to the differences among ENSO phases. The table below reports the results of one-sided Kolmogorov-Smirnov and Shapiro-Wilk tests for normality – even the simulated yield data shows that yield distributions are significantly different from their normal counterparts. The results imply that the yields during none of the three ENSO cycles, and even the pooled yields, are normally distributed with reasonable confidence levels.

K-S test against normality (combined)

	D	P-value	Corrected
All Years	0.0632	0.000	0.000
ElNino	0.0558	0.082	0.073
LaNina	0.0820	0.001	0.001
Neutral	0.0788	0.000	0.000
Shapiro-Wilk W test for normal data			
	W	z	Prob>z
All Years	0.95871	10.509	0.00000
ElNino	0.97961	4.690	0.00000
LaNina	0.96183	6.560	0.00000
Neutral	0.93921	9.929	0.00000

Finally, plots of the simulated yield non-parametric distribution densities are shown in the figure below.

Kernel Densities of Simulated Yields (Gaussian kernels, Silverman's width)

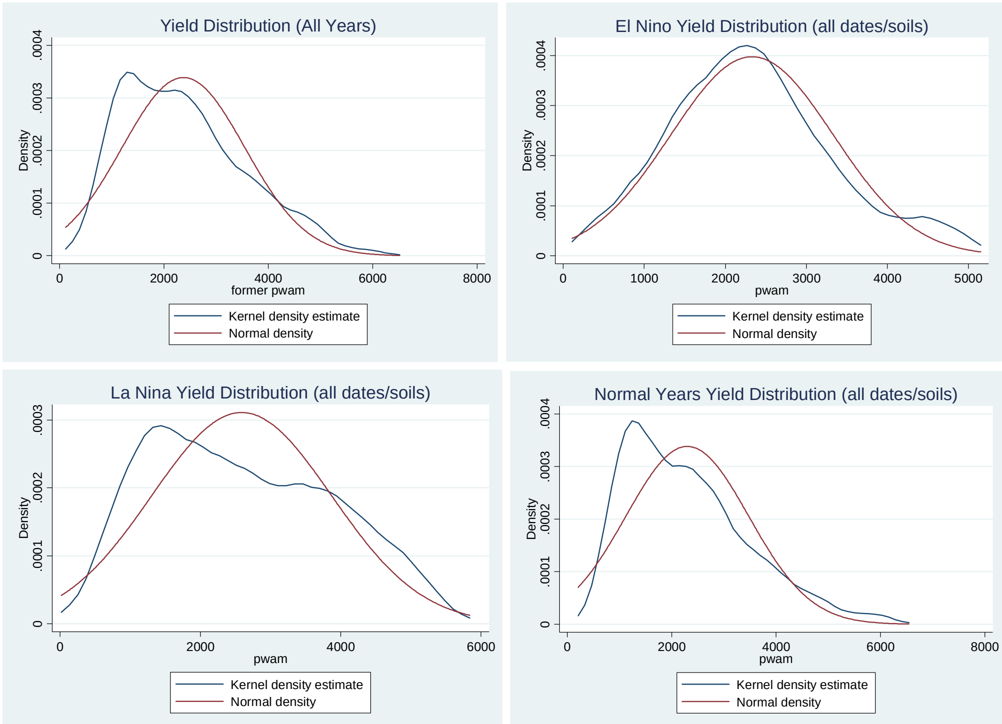

Visual inspection suggests that ElNino yields seem to be more normally distributed than even the pooled data. Perhaps this means that, indeed, the left skewness of yield data is the norm, and the ElNino yields deviate from this general pattern because ElNino years are wetter in the NW Florida and SW Georgia, making below average yields less likely. This might also help explain the lowest kurtosis of the distribution of the dry LaNina yields.

Historical County Yield Data.

Data Description

County-level annual yield series (appropriate for GRP contracts) were used in the analysis. Three crops, corn, peanuts, and cotton were chosen as being most widely produced in the region. The planting areas in the southeastern tri-state region were divided into South-Central Georgia, South Alabama, North-Central Alabama, North-Western Florida on the basis of distinct ENSO impacts. Within the regions, the counties were chosen by the production volumes as defined by the NASS county estimates. As an example, the picture below shows the major corn producing counties in the SE Georgia in 2005. Accordingly, the analysis included all corn producing counties with yields above 10,000 bu.

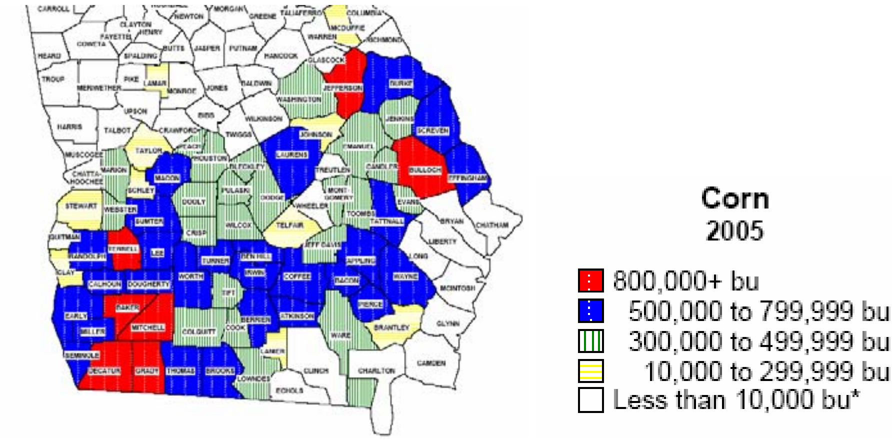

Altogether the set includes 94 corn county yield series, 95 cotton series, and 81 peanut series. The series time span ranges from 78 to 45 years (1934-2005), depending on crop/area, which is just enough to produce non-parametric density estimates. Obviously, the series can not be used in their pristine as they are nonstationary.

Non-Stationarity and Detrending of Historical Data

In order to properly single out the impact of the ENSO phases on yield distributions, it is necessary to remove trends, autoregressive effects, heteroscedasticity, and other systemic influences affecting the distributions. Many of these are expected in the yield time series mainly due to technological changes (trends and heteroscedasticity) and drought/moisture effect persistence (auto-regression and moving average). Below we discuss these and some other problems with the yield data and present the results of testing for them.

Heteroscedasticity:

The Breusch-Pagan-Godfrey test is used to test for heteroscedasticity in the yield series. The table below shows the test statistics for the three crops in the Southeast, which confirm that the raw yield series are definitely heteroscedastic. Cotton yields are slightly more heteroscedastic than corn and peanuts. An interesting observation is that taking logs of the yields reduces but does not eliminate heteroscedasticity (still rejecting H_0). The most commonly suggested reason for heteroscedasticity is the technological progress which, increases yield variance together with the mean.

Heteroscedasticity of yield series

CROP	Breusch-Pagan statistic	Prob > chi2
Corn	6.05	0.0139
Cotton	9.55	0.002
Peanuts	5.45	0.0195

The problem of heteroscedasticity in detrending the yield series can be alleviated by using White's Heteroscedasticity-Consistent Variances and Standard errors OLS estimation. The resulting errors are also known as **robust standard errors** or **robust variances** (parameter estimates are unchanged). An alternative is to calculate Newey-West standard errors for coefficients estimated by OLS regression. The error structure is assumed to be heteroskedastic and possibly autocorrelated up to some lag. There is also the Prais-Winsten and Cochrane-Orcutt regression that accommodates autocorrelation of residuals. However, these techniques only marginally improve on the parameter significance.

Autocorrelation and unit roots:
Autocorrelation in crop yield series has been widely documented. The effects of drought and moisture carryover have been suggested as the main reasons. As autocorrelation in regressors impairs estimates and autocorrelation in the error terms (serial correlation) leads to series non-stationarity, detecting it is important.

While the most standard test for autocorrelation is Durbin-Watson, it has pitfalls due to the regions of indecision. We use the Breusch-Godfrey (BG, also LM) test for autocorrelation the results of which are reported in the table below (1 lag, higher rhos' significance declines slowly):

Serial correlation in the yield series

CROP	Breusch-Godfrey LM test statistic	Prob > chi2
Corn	11.379	0.0007
Cotton	3.515	0.0608
Peanuts	18.077	0.0000

Correlograms also confirm the presence of serial correlation in the yield series, the most significant being in peanuts and only marginally significant in cotton. Apart from the suggested moisture/drought carryover effects and possible technological and behavioural adjustments (usually responsible for serial correlation in economic time series), one might also suspect persistence of ENSO effects but, to the best of our knowledge, ENSO effects persist no longer than a few months.

The correlograms show slowly declining autocorrelation coefficients, confirmed by the Portmanteau Box and Pierce Q-statistic, suggesting a possible RW process. The statistic has chi-squared distribution with df = lag length and is computed for every lag length. If Q exceeds critical value for chosen significance, one rejects the null of all zero rho's – at least some of them must be non-zero (non-stationary time series). The figure below shows two correlograms typical for the corn series – one for the level and the other for first differenced series. The correlogram patterns of autocorrelation and partial autocorrelation (alternating signs) suggest the presence of both autoregression and moving average.

Correlogram of the corn yields in levels

```
                                          -1       0       1    -1       0       1
 LAG      AC       PAC      Q     Prob>Q   [Autocorrelation]  [Partial Autocor]
-------------------------------------------------------------------------------
1       0.9041   0.9651   45.819  0.0000             |-------            |-------
2       0.8638   0.4247   88.469  0.0000             |------             |---
3       0.8239   0.3436   128.04  0.0000             |------             |--
4       0.7747  -0.0047   163.75  0.0000             |------             |
5       0.7402   0.1413   197.02  0.0000             |-----              |-
6       0.6809   0.1315   225.77  0.0000             |-----              |-
7       0.6161   0.0389   249.83  0.0000             |----               |
8       0.5566  -0.0588    269.9  0.0000             |----               |
9       0.5020   0.0776    286.6  0.0000             |----               |
10      0.4568  -0.1514   300.74  0.0000             |---               -|
11      0.4227   0.1658   313.15  0.0000             |---                |-
```

```
12      0.3630  -0.3038   322.51  0.0000            |--                --|
13      0.3162   0.3395    329.8  0.0000            |--                  |--
14      0.2806   0.1101   335.69  0.0000            |--                  |
15      0.2203   0.2645   339.41  0.0000            |-                   |--
16      0.1783   0.3394   341.91  0.0000            |-                   |--
17      0.1021  -0.1221   342.76  0.0000            |                    |
18      0.0535   0.0793      343  0.0000            |                    |
19      0.0014   0.1296      343  0.0000            |                    |-
20     -0.0671  -0.1717   343.39  0.0000            |                   -|
21     -0.1101   0.1593    344.5  0.0000            |                    |-
```

<u>Correlogram of differenced corn yield series</u>

```
                                              -1      0      1   -1      0      1
 LAG      AC      PAC      Q     Prob>Q    [Autocorrelation]  [Partial Autocor]
-------------------------------------------------------------------------------
1      -0.4087  -0.4324   9.1961  0.0024         ---|             ---|
2      -0.1048  -0.3451   9.8126  0.0074            |              --|
3       0.2031   0.0038   12.177  0.0068            |-               |
4      -0.1671  -0.1411   13.811  0.0079           -|               -|
5      -0.0308  -0.1299   13.868  0.0165            |               -|
6       0.0830  -0.0395   14.289  0.0266            |                |
7       0.0065   0.0576   14.291  0.0462            |                |
8      -0.0962  -0.0794   14.882  0.0615            |                |
9       0.1693   0.1487   16.754  0.0527            |-               |-
10     -0.1575  -0.1690   18.413  0.0484           -|               -|
11      0.1837   0.2999   20.725  0.0363            |-               |--
12     -0.2254  -0.3442   24.293  0.0186           -|              --|
13      0.0437  -0.1046   24.431  0.0274            |                |
14      0.0874  -0.2618   24.995  0.0346            |              --|
15     -0.2284  -0.3426   28.953  0.0163           -|              --|
16      0.2451   0.1000    33.64  0.0061            |-               |
17     -0.1314  -0.1052   35.025  0.0062           -|                |
18     -0.0230  -0.1891   35.068  0.0093            |               -|
19      0.1217   0.0261   36.329  0.0096            |                |
20     -0.0697  -0.3018   36.755  0.0125            |              --|
21     -0.1121  -0.3746   37.893  0.0133                           --|
```

The analysis suggests that the series, at least in levels, may have a unit root as a reason for its non-stationarity. The table below reports the results of the Dickey-Fuller and the Phillips-Perron unit root tests. The latter is preferable because it takes care of possible serial correlation in the error terms by adding lagged difference terms of the regressand. The asymptotic distribution is the same as for the DF. The tests were performed on both the raw yield series for each crop and on the linearly detrended (with robust variance estimates) yield series. This is important in considering difference versus trend stationarity of the series.

Unit-root test results, level series

CROP	Dickey-Fuller yield series statistic	DF linearly detrended series statistic	5% critical value	Phillips-Perron Statistic	PP linearly detrended series statistic	5% critical value
Corn	-0.800	-4.356	-2.928	-0.446	-29.317	-13.316
Cotton	-1.826	-4.218	-2.920	-4.222	-30.283	-13.396
Peanuts	-1.583	-2.813	-2.913	-3.002	-12.971	-13.468

The detrended series in levels are likely to possess unit roots; however, many linearly detrended series do not have unit root, particularly in case of corn. This suggests yields might be trend stationary processes (TSP) instead of difference stationary processes (DSP). The table below shows the results of unit root testing of the first differenced yield series.

Unit-root test results, first difference series

CROP	Dickey-Fuller yield series statistic	5% critical value	Phillips-Perron Statistic	5% critical value
Corn	-10.746	-2.929	-64.336	-13.308
Cotton	-12.336	-2.923	-78.192	-13.372
Peanuts	-11.610	-2.914	-11.970	-2.914

However, as certain variation in unit root test results exists among the county yield series (the reported values are just averages), we did not automatically assume TSP over DSP. Instead, we tried both linear detrending of the series using robust variance estimates and various specifications of the ARIMA models, as well as ARIMA(p, 0, q) run on linearly detrended series.

The model selection criteria were the (1) Akaike and (2) Schwarz's Bayesian information criteria. Both are better than the R^2 not only in that they punish for adding regressors but also because they are good for out of sample forecasts. The AIC is also used to determine the lag length in the AR(p) models. Another model selection criterion was the residuals being white noise (using correlograms and the tests described above).

<u>Cointegration in yield series</u> (among similar ones):
To confirm consistency of the historical county average yield series, the Breusch-Godfrey (LM), Engle-Granger, and cointegrating regression Durbin-Watson (CRDW) tests were performed. The procedures test the residuals from regressing one time series on another for unit root and autocorrelation, also confirmed by correlogram analysis. As expected, same crop and geographical area yield series were found to be highly cointegrated, whereas different crop and area yields were only slightly less cointegrated. Also, the residuals from regressing even "unrelated" series showed little signs of heteroscedasticity, based on the Breusch-Pagan / Cook-Weisberg test. This may suggest that the yields were largely influenced by the climate variability. The table below shows some typical results of the cointegration testing.

Cointegration tests

Cointegration	R-squared	Dickey-Fuller Z(t)	Phillips-Perron Z(t)	Breusch-Godfrey Prob>Chi2	Breusch-Pagan Prob>Chi2
Corn-Corn	0.65	-5.946	-5.937	0.7730	0.8795
Cotton-Cotton	0.8	-6.296	-6.451	0.2661	0.3302
Peanuts-Peanuts	0.91	-5.448	-5.555	0.1855 0.0257	0.2590
Corn-Cotton	0.17	-5.126	-5.128	0.8551	0.5643

Corn-Peanuts	0.34	-4.335	-4.395	0.0473	0.9157
Cotton-Peanuts	0.28	-4.392	-4.407	0.0001	0.1858

Structural breaks:

Another potential problem is the presence of structural breaks in the yield series, the reason for which might be the introductions of technological innovations (irrigation, etc.). In looking for the structural breaks, it is useful to use recursive LS: start with a sub-sample and add one datum/year at a time re-estimating the model (rolling window). The search for structural breaks in our dataset did not reveal any (also confirmed by the Chow's Prediction Failure test that compares the actual out-of sample data with predictions).

Volatility clustering**:**

Volatility clustering (heterogeneity and autocorrelation in the error variance) has been observed in economic series (particularly in the stock market indices), and plausible behavioral and technological explanations for it have been suggested. While there is no a priori reason to expect volatility clustering in yield series, we estimated autoregressive conditional heteroscedasticity (ARCH) models, for which squared residuals from a regression of yield series on time (squared deviations from the mean of detrended series) were regressed on own lagged values. The results showed highly insignificant coefficients at the lagged values, leading to a conclusion that the series do not possess autocorrelation in the error variance (hence no volatility clustering).

Taking logs of the yield series:

Taking logs of time series has been suggested as a universal remedy for serial correlation and autoregression. This seems to be true in case of the southeastern crop yields:

- Phillips-Peeron and Dickey-Fuller tests of the logged series show only weak evidence of unit roots, and usually no unit root in the linearly detrended logs of the series. The conclusion is supported by correlogram analysis.
- The Breusch-Godfrey LM test for autocorrelation shows evidence of zero autocorrelation in the error terms. The only reason to expect this is that, as crops are planted in rotation, there're no moisture carryover effects.
- Breusch-Pagan-Godfrey / Cook-Weisberg test shows no heteroscedasticity in the logged yields.

However, as is shown in the results section below, yield distribution estimation based on linearly detrended logged yield series (take logs, detrend, unlog, proceed density estimation) shows results consistent with both linear detrending of raw yield series and ARIMA residuals. The only difference is that the actuarially fair premiums estimated using logged series are a little higher – probably the result of the transformation. Overall, the results are pretty robust to the different models used in detrending and "stationarizing" the series.

ARIMA estimation:

Estimation of correctly specified ARIMA (autoregressive integrated moving average) models removes the non-stationarity of time series that is due to autoregression and serial correlation. A simple ARIMA(p,d,q) model lets the data explain itself by regressing a variable on its own differenced, lagged, and error values:

$$Y_t^* = \varphi_1 Y_{t-1}^* + \varphi_2 Y_{t-2}^* + ... + \varphi_p Y_{t-p}^* + \varepsilon_t + \theta_1 \varepsilon_{t-1} + ... + \theta_q \varepsilon_{t-q}$$

where Y^* is d-times differenced value of the series and the ε's are IID normal errors. Parameters p and q specify the order of AR and MA processes, and d specifies the degree of differencing the analyzed variable that is necessary in case of unit roots.

Using the Akaike and Schwartz information criteria, we found ARIMA(1,1,2) specification to be the most adequate for most of the yield series. In reporting the results below, unless otherwise specified, we use distributions derived from the ARIMA(1,1,2) residuals. However, the results seem to be robust to detrending/refining methodology, and most often yield distribution estimates derived from linearly detrended series and exponential smoothing (a simplified version of ARIMA) show the same patterns of the ENSO phase dependent actuarially fair premiums.

Comparison of Simulated and Empirical Yield Distributions

In comparing the historical county-level yields with the simulated ones, we did not expect to get a perfect match. In particular, there is no reason to expect the average county yields to be similar as the NASS data includes both irrigated and non-irrigated yields (for instance, about 20% of the peanut acreage in the southwest Georgia has been irrigated during the last 20 years, and the yields from irrigated production are typically ~30% higher). Irrigated yields are less volatile due to independence of precipitation, which dampens any possible ENSO influences. The ENSO effects are further reduced because of the averaging of individual data in the county estimates. However, there is no reason not to expect similarities between the distributions of the real and simulated yields as those are shaped by a group of biological and physiological factors common to both practices. The table below lists the four moments of the Colquitt county empirical peanut yield distributions.

Colquitt county (GA) peanut yield distribution moments

enso	**mean**	**sd**	**skewness**	**kurtosis**
ElNino	3316.9	362.5	-0.073	2.77
LaNina	3369.8	478.8	0.441	1.838
Neutral	3192.6	555.2	0.178	2.735
All years	3260.5	503.8	0.110	2.802

The average yields by ENSO phase preserve the same relation as the simulated data: LaNina > ElNino > Neutral. However, he average yields are statistically different only at levels greater than 23.4%. While the reasons for LaNina yields being the largest are still unclear, the Neutral year yields are probably the smallest because of the suggested higher freeze probability. The relative magnitudes of variance, skewness, and kurtosis seem to be different from the simulated yields, but then their ENSO phase differences are not statistically significant. The significantly smaller variance of actual yields is explained by irrigation and by county-level averaging. Normally, higher differences between simulated and county yield variances should indicate greater heterogeneity among individual producers, which may account for the differences in the other distribution parameters (Barnett, Black, Hu, Skees, 2005).

As was mentioned before, higher moments are less relevant when dealing with non-parametric distributions, when visual inspection is of greater importance. The table below presents comparisons of the kernel densities of the simulated and actual yield distributions by ENSO phase with overlaid normal densities (in red).

Simulated and Historical Yield Distributions, Colquitt, GA (normal overlay)

All Years	ElNino	LaNina	Neutral
Colquitt County Simulated yields (all planting dates, all soils)			
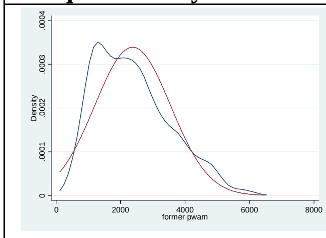	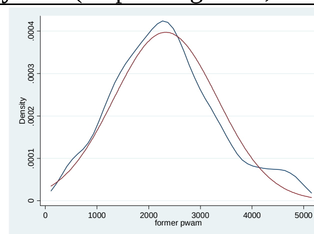	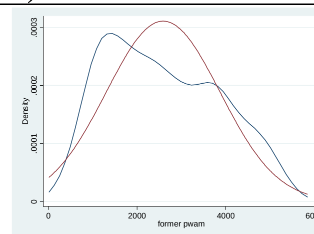	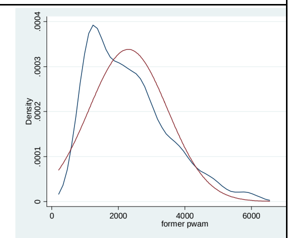
Colquitt County Simulated yields (optimal planting dates, all soils)			
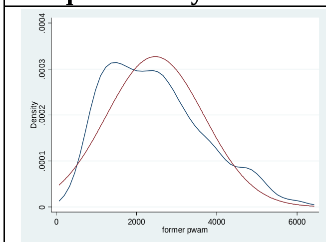	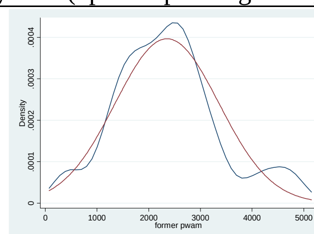	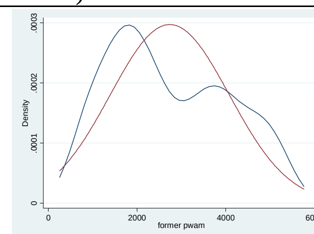	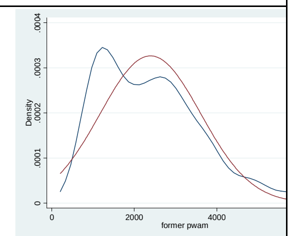
Colquitt County Average Yields, 1934-2005			
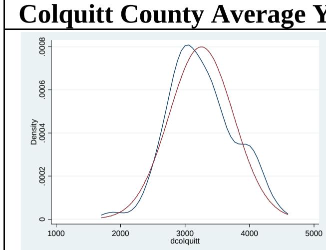	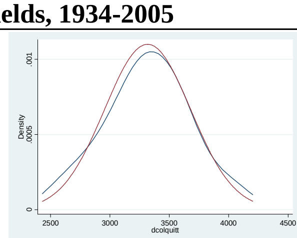	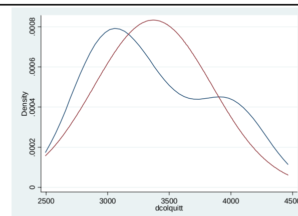	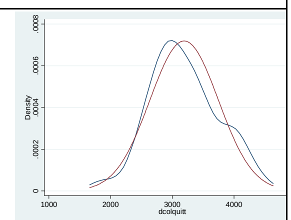

The ENSO-dependent distribution similarities are largely preserved across counties. In both simulated and empirical series, the ElNino densities seem the least skewed and closest to normal, while the LaNina densities show the greatest evidence of bi-modality. The Neutral year densities are the highest peaked, as evidenced by the kurtosis value. In most cases, LaNina yields also show the strongest, and Neutral year the weakest, evidence of bi-modality which is also reflected in their variance. In drawing these conclusions, it should be mentioned that the data span is sufficiently long making it unlikely that the differences in the ENSO dependent distribution patterns are due to chance.

For illustrative purposes, we computed the frequencies of the incidences of yield distribution moments being the highest or the lowest during an ENSO phase for the 17 counties and compared them to the simulated data:

Frequencies of *Empirical* Distribution Moments being Highest and Lowest among the ENSO Phases, % of total

	mean			std			skewness			kurtosis		
	EN	LA	NE	EN	LA	NE	EN	LA	NE	EN	LA	NE
max	29	**71**	0	0	**82**	18	29	**65**	6	12	24	**65**
min	0	6	**94**	**100**	0	0	47	6	47	47	47	6
Frequencies of *Simulated* Distribution Moments being Highest and Lowest among the ENSO Phases, %												
	mean			std			skewness			kurtosis		
	EN	LA	NE	EN	LA	NE	EN	LA	NE	EN	LA	NE
max	0	**100**	0	0	**100**	0	0	0	**100**	0	0	**100**
min	0	0	**100**	**100**	0	0	0	**100**	0	0	**100**	0

Clearly, there are similarities. The empirical data show that the majority of highest county average yields happen in the LaNina years and the overwhelming majority of lowest average yields happen in the Neutral years, which agrees with the simulated data. The empirical data also indicate that about 30% of the highest average county yields happen in the ElNino years. The high yields in LaNina years may be due to the observed higher cumulative solar radiation and the low yields during Neutral years may be due to the lack of it and also to the relative lack of rainfall. It is reasonable to expect that, *ceteris paribus*, higher yields during LaNina years should reduce expected insurance losses. The standard deviations of the average county yield series are the highest during LaNina in 82 percent of the counties and they are always the lowest during ElNino, which also matches the simulated data. However, the skewness coefficient's relative frequency does not match the simulated data. Finally, the kurtosis is the highest during the neutral years in 65% of the counties analyzed and happen with equal frequencies during ElNino and LaNina years.

Apart from this, it is hard to describe these empirical distributions in much greater detail. The significance of distributional differences is determined by their effect on the application that uses these data. Our primary interest is the area yield insurance implications of using non-parametric distribution estimates.

Comparison of Insurance Premiums

The methodology for calculating ENSO phase dependent insurance premiums is described in the Methodology section. The omission of unusually low (catastrophic) yields in the non-parametric densities/distributions drives down the estimated premium rates. Without reliable information on how ENSO phases affect crop failures, there is little reason to try to accommodate them here. Thus, comparison with the actual GRP premiums is premature at this point, and the term "premiums" is used interchangeably with "expected losses".

Below are two tables showing expected losses for Colquitt county data.

GRP Insurance Rates by ENSO Phase Derived from *Simulated* Colquitt County Yields

Actual Ranges			Expected Loss to Coverage Ratio				
MEAN	2509.886		All Years	ElNino	LaNina	Neutral	Normal
COVERAGE	70%	1756.92	8.51%	5.99%	8.30%	9.63%	6.59%
	80%	2007.909	11.39%	8.05%	10.71%	12.79%	8.89%
	90%	2258.897	14.49%	10.52%	13.29%	16.05%	11.51%

GRP Insurance Rates by ENSO Phase Derived from *Historical* Colquitt County Yields

Actual Ranges (smaller rates)			**Expected Loss to Coverage Ratio**				
Av. Yield	3260.508		All years	ElNino	LaNina	Neutral	Normal
COVERAGE	80%	2608.406	0.63%	0.01%	0.07%	1.21%	0.75%
	90%	2934.457	2.31%	0.92%	0.99%	3.51%	2.44%

Legend: __ - largest values; __ - smallest values.

The big difference between the simulated and actual data rates is due to (1) the averaged nature of the county data (only systemic risk), (2) the presence of much more stable irrigated yields in the county data, and (3) the broader range of simulated (individual) yields.

These differences notwithstanding, in this research, we are interested exclusively in the ordinal properties of the expected losses, i.e., in their differences between ENSO cycles. There is clearly a pattern in the differences among the expected losses that persists in both the simulated and the empirical data. In both cases, premiums (losses) are clearly the highest in the Neutral years and the lowest in the ElNino years for all coverage levels. This is a result of the generally lower mean and higher kurtosis in ElNino years, meaning that the bulk of the mass is immediately to the left of the series average, hence the higher expected loss. Predictably, the expected losses estimated with pooled data fall in between the extremes. It is also worth noting that the expected losses based on the assumption of normality (using the series' mean and std) are relatively higher for the actual data and for the simulated data, which is probably a result of a smaller span of the former.

The premiums (losses) are also bigger in LaNina than in ElNino years. This is because the ElNino distribution is less skewed (although the variance seems counter-intuitive). Obviously, estimated loss probabilities and premium rates approach zero in spite of the fact that the true expected losses are likely to be positive, albeit small. Catastrophic loss probability can be constructed to address this issue (see Goodwin and Ker). However, investigation of the ENSO impact on catastrophic crop losses is beyond the scope of this paper.

The differences in the absolute values of expected losses among the counties are due to different variance of the county series (ranges from 387 to about 650), which is a universal factor for "inter-county" estimates (ENSO phases for the same county) but differs between the counties.

Again, it is important to note that these results are not readily comparable to the actual premiums because the latter are not likely to be actuarially fair and because our data do not include catastrophic losses (usually defined as yields more than two standard deviations below last four years' average). The most important finding is that they differ among the ENSO phases, and that can only be explained by the differences in their true non-parametric (as opposed to assumed theoretical) densities, even if the moments of the distributions seem to be similar.

The simulated and historical yield distribution differences notwithstanding, the analysis of historical yield data on corn, cotton, and peanuts reveals distinct ENSO dependent patterns in non-parametrically estimated insurance rates across counties and regions in the Southeastern US reported in the next section.

Insurance Rates Based on Historical Yield Data

In this section, we report the results of actuarially fair GRP rate estimation using county average annual yield time series. As was mentioned before, the rates are calculated for yield distributions for distinct ENSO phases (ElNino, LaNina, and Neutral). The distributions are estimated non-parametrically using kernel density estimates. The ENSO data used for grouping the yield series was constructed specifically for the purpose by adjusting the monthly ENSO indices to reflect the ENSO conditions prevailing during the crops' growth season, not calendar time.

The crops used in the analysis are corn, cotton, and peanuts. The areas from which the data were collected include South-Central Georgia, South Alabama, North-Central Alabama, and North-Western Florida (defined by major producing counties as reported by the NASS). The time span for the series ranges form 78 to 45 years, which includes 16-11 ElNino, 17-10 LaNina, and 45-25 Neutral years, depending on crop/area.

Alternative methodologies used for detrending and "stationarizing" the yield series include:

- Linear using robust variance estimator;
- Double exponential smoothing (used in crop insurance);

- ARIMA models, parameters determined by testing.

On the basis of statistical tests and model fitting criteria, ARIMA(1,1,2) specification was chosen as most appropriate.

The results of **parametric** distribution analysis and testing of yield data thus refined can be summarized as follows:

- There are no ENSO-dependent differences in yield averages or variances;
- Normality tests are rejected in most cases;
- There are some ENSO-dependent patterns in yield distribution skewness (discussed above);
- Most importantly, GRP rates calculated using parametric yield distribution specification show no ENSO impact.

However, yield distribution densities derived using **non-parametric** kernel density estimation do show ENSO-dependent differences in most cases. While visual inspection of detrended yield density plots reveals noticeable but not easily quantifiable differences, the insurance rates calculated using the above methodology show distinct ENSO-phase dependent patterns. Graphically, the area yield insurance premiums calculated using Gaussian kernel density estimation of ARIMA(1,1,2) detrended county yield series compare as follows (averages of county groups by crop and region).

<u>Corn</u> yield loss to coverage rations (actuarially fair premiums) appear to be the most dependent on the ENSO phases, the premiums being the highest, regardless of coverage, during ElNino and the lowest during LaNina years (except for Southern AL). The figure below shows plots of the rates for the four geographical regions.

GA Corn, South-Central

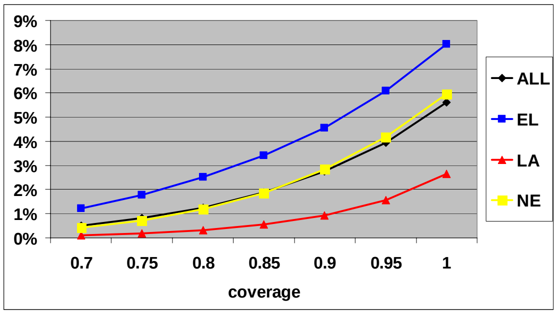

AL Corn, North-Central

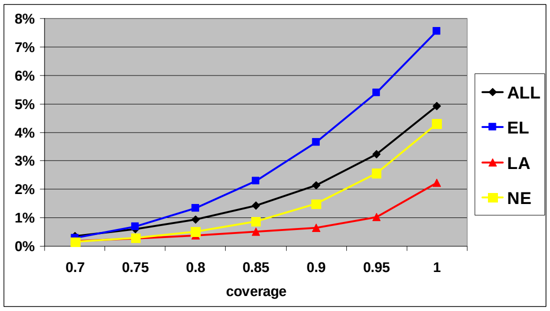

AL Corn, South

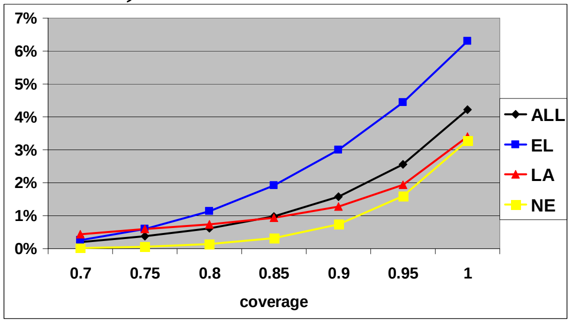

FL Corn, North-West

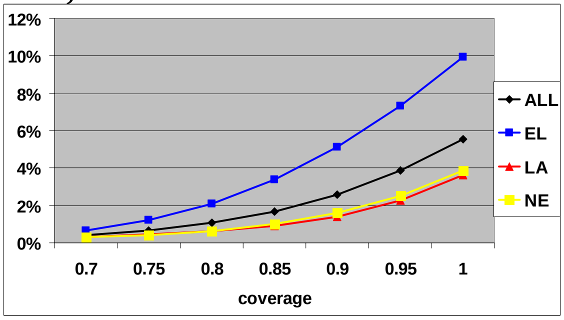

These results appear to be robust to the detrending methodologies tried in the analysis. Comparison of the rate schedules in the figure below shows that, while the magnitudes of the

ENSO phase dependent insurance rates differ, their relative values are preserved throughout the alternative treatments.

GA Corn, Linear Detrending

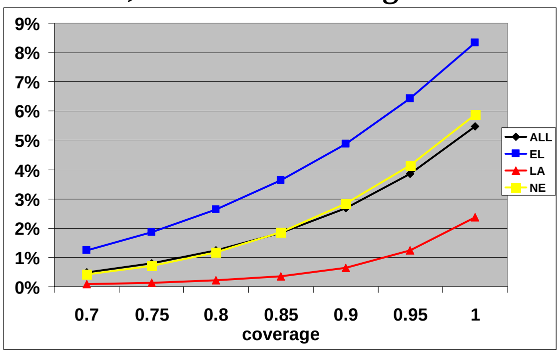

GA Corn, Linear Detrending, Logged Series

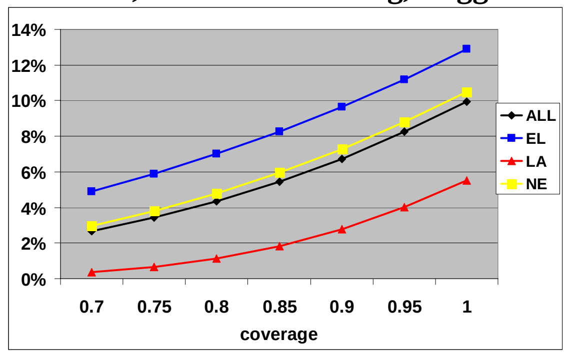

GA Corn, Double Exponential Smoothing

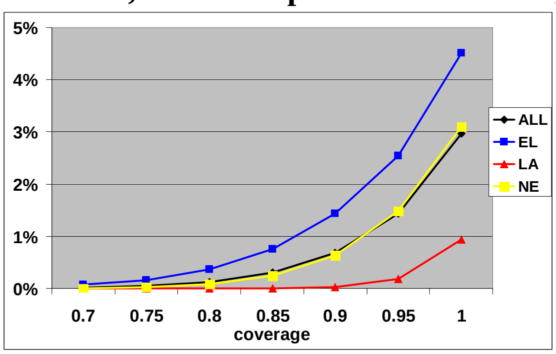

GA Corn, ARIMA(1,1,1)

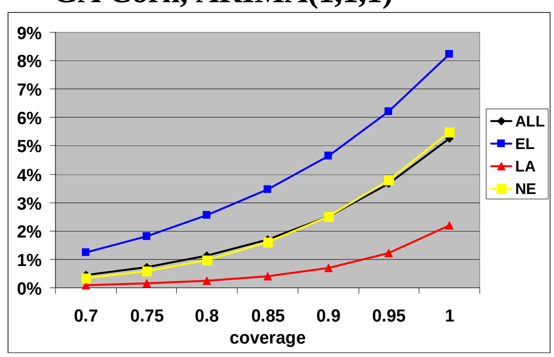

Cotton yield premiums show a pattern similar to that of corn, being the highest during ElNino (except for North-Central AL) and the lowest during LaNina years.

GA Cotton, North-Central

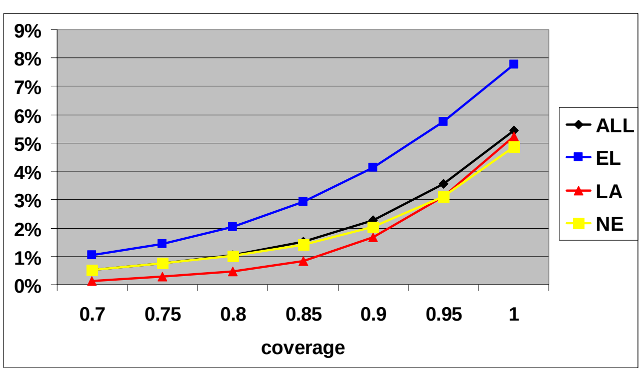

FL Cotton, North-West

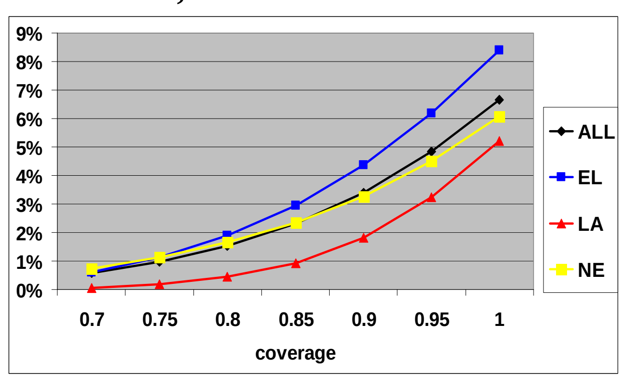

AL Cotton, South

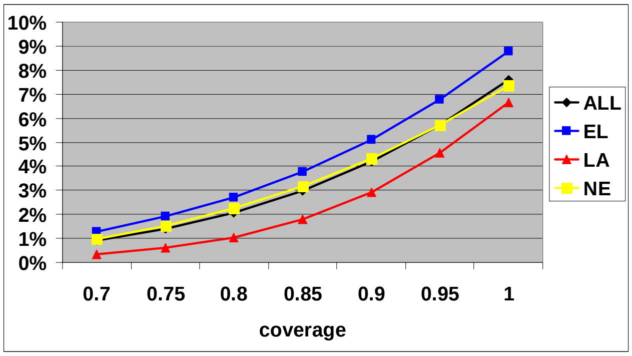

AL Cotton, North-Central

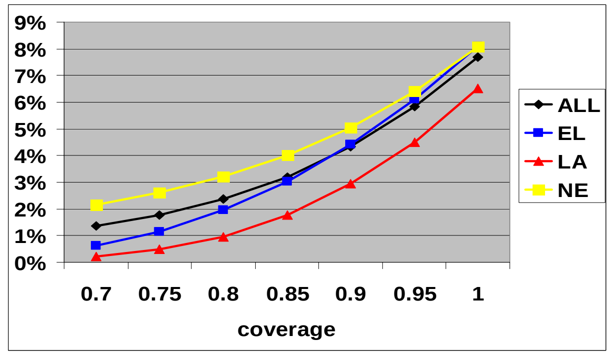

Finally, peanut premiums are higher during Neutral years and lowest during LaNina (except for Southern AL), the differences being muffled in NW Florida.

GA Peanuts, South-Central

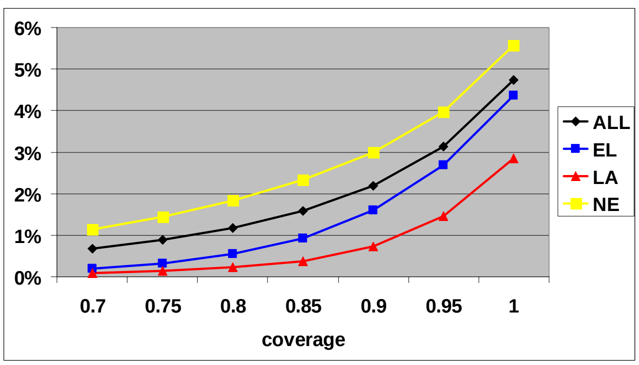

AL Peanuts, South

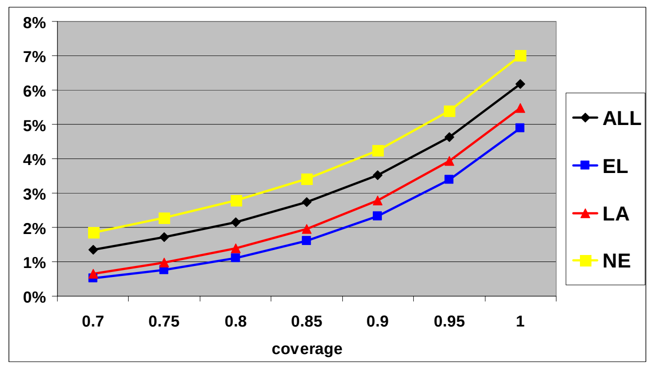

FL Peanuts, North-West

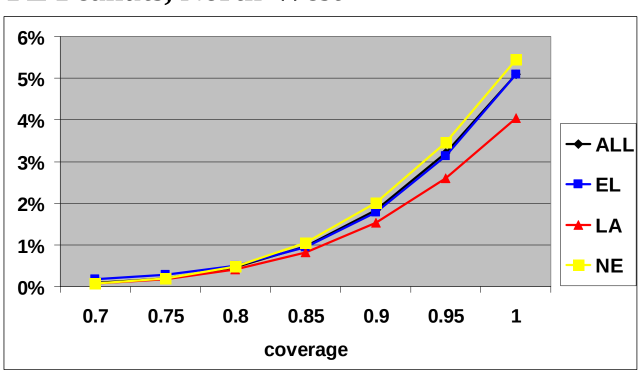

At this point, it is hard to reason how exactly the ENSO climate differences affect yields. However, the estimations for peanuts in GA using historical data correspond to those using peanut growth simulation data in Colquitt county, GA, which lends validity to our findings.

These results suggest that incorporating information on the ENSO cycle in insurance premiums in the Southeast can be beneficial for both producers and insurers, even if the actual insurance premiums include other components. The benefits to the producer depend on the actual risk aversion of the farmer – more accurate estimation of expected losses results in larger utility gains from purchasing insurance that hedges against yield risk. The magnitude of the benefits depends on risk aversion and the amount of insurance purchased, among other factors. Better structured contracts may also increase demand for insurance. The benefits to the insurance companies consist of increased demand for insurance and of avoiding short-term losses because of their more accurate estimation which, depending on the companies' financing policies and planning horizon, might also increase the long-term return on investment.

Conclusions

In this paper, we report results of non-parametric analysis of peanut, corn, and cotton yield distributions by the ENSO phases in the Southwestern Georgia, Northwestern Florida, and Southern and Northwestern Alabama. For comparison and validation purposes, the historical yield data is complemented by a set of simulated peanut yields generated using daily weather, soil, and management practices data in the Colquitt county, Georgia.

The focus of the analysis is on establishing ENSO-dependent differences in the yield distributions and on evaluating their implications for area yield crop insurance the expected losses for which are calculated using county average yield series. The hypothesis is that different climate conditions during ENSO cycles translate into different yield distributions, which is justified by the observed South-Eastern climate differences and previous research showing the importance of ENSO cycles for optimal planting dates.

The results of kernel density estimates of historical county yield data show consistent patterns in actuarially fair insurance rate schedules grouped by ENSO phases, crops, and geographical areas. In particular, corn yield insurance premiums appear to be the most dependent on the ENSO and are the highest, regardless of coverage, during ElNino and the lowest during LaNina. Peanut premiums are higher during Neutral years and lower during LaNina (except of Southern Alabama). Cotton premiums are the highest for ElNino (except for North-Central AL) and the lowest during LaNina years. The results appear to be robust to the transformations used to make the yield series stationary. While these dependencies do not necessarily correspond to the precipitation and solar radiation characteristics of the corresponding ENSO cycles in the Southeastern US, drawing direct analogies with yield variability is premature as many less documented factors, like the spacing of sunny and rainy days, may be just as important.

The validity of these findings is reinforced by comparison of the empirical and simulated peanut yield distributions. The comparisons show that the distributions are similar in many ways and that the dissimilarities can be explained by known factors. These findings should be more relevant for the area yield insurance as opposed to the APH arrangements as the yield data used in designing contracts for the former reflects the systemic risk more dependent on climate than on the farm-level, basis risk factors accommodated in the APH plans.

References:

Barnett, B.J., J.R. Black, Y. Hu, and J.R. Skees. 2005. "Is Area Yield Insurance Competitive With Farm Yield Insurance?" *Journal of Agricultural and Resource Economics*, 30(2): 285-301.

Boote, K.J., J.W. Jones, and G. Hoogenboom. 1998. "Simulation of Crop Growth: CROPGRO Model," *Agricultural Systems Modeling and Simulation*, R.M. PEart and R.B. Curry, Eds., Marcel Dekker, New York, 113-133.

Cabrera, V.E., C.W.Fraisse, D. Letson, G. Podesta, J. Novak. "Impact of Climate Information on Reducing Farm Risk by Optimizing Crop Insurance Strategy," Working Paper, University of Miami.

Fraisse, C.W., A. Garcia y Garcia, J.L. Novak, J.W. Jones, and G. Hoogenboom. "Using Crop Models and ENSO-Based Climate Forecast to Aid in Peanut Crop Insurance Decisions," Working Paper, University of Florida.

Garcia y Garcia, A., G. Hoogenboom, L.C. Guerra, J.O. Paz, and C.W. Fraisse. "Analyzing Long-Term Historical Peanut Yield in Georgia with a Crop Simulation Model," Working Paper, University of Georgia.

Goodwin, B.K., and A.P. Ker. 1998. "Nonparametric Estimation of Crop Yield Distributions: Implications for Rating Group-Risk Crop Insurance Contracts," *American Journal of Agricultural Economics*, 80:139-153.

Hansen, J.W., A.W. Hodges, and J.W. Jones. 1998. "ENSO Influences on Agriculture in the Southeastern United States," *Journal of Climate*, 11:404-411.

Jones, J.W., G. Hoogenboom, C.H. Porter, K.J. Boote, W.D. Batchelor, L.A. Hunt, P.W. Wilkens, U.Singh, A.J. Gijsman, and J.T. Ritchie. 2003. "DSSAT Cropping System Model," *European Journal of Agronomy*, 18:235-265.

Li, Q., J.S. Racine. *Nonparametric econometrics*. Princeton University Press, 2007.

Mavromatis, T., S.S. Jagtap, and J.W. Jones. 2002. "ElNino Southern Oscillation Effects on Peanut Yield and Nitrogen Leaching," *Climate Research*, 22:129-140.

Mjelde, J.W., H.S.J. Hill, and J.F. Griffiths. 1998. "A Review of Current Evidence on Climate Forecasts and Their Economic Effects in Agriculture," *American Journal of Agricultural Economics*, 80(5):1089-1085.

Orlove, B.S., J.C.H. Chiang, and M.A. Cane. 2000. "Forecasting Andean rainfall and crop yield from the influence of El Niño on Pleiades visibility," *Nature* 403, 68-71.

Podesta, G., D. Letson, C. Messina, F. Royce, R.A. Ferreyra, J.W. Jones, J.W. Hansen, I. Llovet., M. Grondona, and J. O'Brien. 2002. "Use of ENSO-related Climate Information in Agricultural Decision Making in Argentina: A Pilot Experience," *Agricultural Systems*, 74(3):371-392.

Sherrick, B.J., F.C. Zanini, G.D. Schnitkey, and S.H. Irwin. 2004. "Crop Insurance Valuation under Alternative Yield Distributions," *American Journal of Agricultural Economics*, 86(2):406-419.

Silverman, B.W. *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, New York, 1986.

Skees, J.R., J.R. Black, and B.J. Barnett. 1997. "Designing and Rating an Area Yield Crop Insurance Contract," *American Journal of Agricultural Economics*, 79(2):430-438.